



- (51) International Patent Classification:  
G06F 17/30 (2006.01)
- (21) International Application Number:  
PCT/CN2016/082980
- (22) International Filing Date:  
23 May 2016 (23.05.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; One Microsoft Way, Redmond, WA 98052 (US).
- (72) Inventors; and
- (71) Applicants (for US only): BAI, Jing [US/US]; Microsoft Corporation One Microsoft Way, Redmond, WA 98052 (US). LIU, Yue-Sheng [US/US]; Microsoft Corporation One Microsoft Way, Redmond, WA 98052 (US). PEDERSEN, Jan O. [US/US]; Microsoft Corporation One Microsoft Way, Redmond, WA 98052 (US). YANG, Mao [CN/CN]; Microsoft Corporation One Microsoft Way, Red-

mond, WA 98052 (US). LU, Qi [US/US]; Microsoft Corporation One Microsoft Way, Redmond, WA 98052 (US).

(74) Agent: SHANGHAI PATENT & TRADEMARK LAW OFFICE, LLC; 435 Guiping Road, Shanghai 200233 (CN).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,

(54) Title: RELEVANT PASSAGE RETRIEVAL SYSTEM




FIG. 1

(57) Abstract: A new architecture is provided to support a precise information retrieval system on a web scale. The architecture provides algorithms to generate candidates and select the top N results via ranking models (e.g., Semantic ranking models, Aggregation ranking models) to capture term relationships between query and result contents at search-time.

WO 2017/201647 A1

EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,  
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,  
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,  
KM, ML, MR, NE, SN, TD, TG).

**Published:**

— *with international search report (Art. 21(3))*

## RELEVANT PASSAGE RETRIEVAL SYSTEM

### BACKGROUND

**[0001]** Users of computing systems run searches for electronic objects using queries. For example, users run searches on the Internet, email systems, video archives, and other databases. As the number of electronic objects being stored grows, it is becoming increasingly challenging to search a corpus of electronic objects and return relevant results to a user. Further, users increasingly expect to quickly access information relevant to the query without having to access the various electronic objects that may be returned by the query. This is particularly true for queries that are submitted using a mobile device. The small form factor of mobile devices makes it more difficult for users to sift through the electronic objects returned in response to a query. Because of this, it is beneficial to provide a mechanism in which an answer to a query is provided directly to the user without requiring the user to actually access an electronic object to find the answer.

**[0002]** It is with respect to these and other general considerations that the aspects disclosed herein have been made. Also, although relatively specific problems may be discussed, it should be understood that the examples should not be limited to solving the specific problems identified in the background or elsewhere in this disclosure.

### SUMMARY

**[0003]** This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description section. This summary is not intended to identify all key or essential features of the claimed subject matter, nor is it intended as an aid in determining the scope of the claimed subject matter.

**[0004]** Aspects of the present technology relate to returning results from user queries. For example, a user may provide a query, such as a natural language query, seeking information into a web browser, an email search interface, or cloud search interface, a file system search interface, or any other type of search interface. Aspects of the technology described herein provide systems and methods to identify highly relevant passages from a corpus of electronic objects (such as the web pages, word processing documents, spreadsheets, videos, etc.) and return the most relevant passage(s) that answers the user query. The passage may be obtained directly from the electronic object or may be generated from multiple, highly-ranked passages from one or more electronic objects.

Other information may be returned such as the location of the electronic object from where the passage(s) was obtained. This may be a URL link, an email link, or other object link.

**[0005]** Examples are implemented as a computer process, a computing system, or as an article of manufacture such as a device, computer program product, or computer readable medium. According to an aspect, the computer program product is a computer storage medium readable by a computer system and encoding a computer program comprising instructions for executing a computer process.

**[0006]** The details of one or more aspects are set forth in the accompanying drawings and description below. Other features and advantages will be apparent from a reading of the following detailed description and a review of the associated drawings. It is to be understood that the following detailed description is explanatory only and is not restrictive of the claims.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

**[0007]** The accompanying drawings, which are incorporated in and constitute a part of this disclosure, illustrate various aspects. In the drawings:

**[0008]** Figure 1 illustrates an exemplary networked-computing environment for retrieving relevant passages from a corpus of electronic objects.

**[0009]** Figure 2 illustrates an exemplary method for providing a passage.

**[0010]** Figure 3 illustrates an exemplary passage retrieval system for returning relevant passages based on a query.

**[0011]** Figure 4 illustrates an example output of a relevant passage and a link to an exemplar electronic object.

**[0012]** Figure 5 illustrates an exemplary a method for identifying a query dependent passage.

**[0013]** Figure 6 illustrates a method for identifying a query independent passage in a document.

**[0014]** Figure 7 illustrates an exemplary method for ranking a passage based on a semantic translation model.

**[0015]** Figure 8 illustrates an example output resulting from identifying features of a passage based on a query and a sample target electronic object.

**[0016]** Figure 9 illustrates an exemplary method of matching a passage type to a query type.

**[0017]** Figure 10 illustrates an example output generated using a method for matching a passage type to a query type.

**[0018]** Figure 11 illustrates an exemplary method of matching an electronic object to a query based on contextual meaning.

**[0019]** Figure 12 provides an example of matching an electronic object to a query based on a heading of the electronic object.

**[0020]** Figure 13 illustrates an exemplary method of aggregation model of ranking passages.

**[0021]** Figure 14 is an example of applying the aggregation method on the passage.

**[0022]** Figure 15 illustrates an exemplary method of matching answer patterns in passages to query answer patterns for queries.

**[0023]** Figure 16 provides an example of the results generated using a method of matching answer patterns to queries.

**[0024]** Figure 17 is a block diagram illustrating example physical components of a computing device with which aspects of the disclosure may be practiced.

**[0025]** Figure 18A and 18B are simplified block diagrams of a mobile computing device with which aspects of the present disclosure may be practiced.

**[0026]** Figure 19 illustrates one example of the architecture of a system for web-scale passage retrieval as described above.

#### **DETAILED DESCRIPTION**

**[0027]** Various aspects of the disclosure are described more fully below with reference to the accompanying drawings, which form a part hereof, and which show specific exemplary aspects. However, different aspects of the disclosure may be implemented in many different forms and should not be construed as limited to the aspects set forth herein; rather, these aspects are provided so that this disclosure will be thorough and complete, and will fully convey the scope of the aspects to those skilled in the art. Aspects may be practiced as methods, systems or devices. Accordingly, aspects may take the form of a hardware implementation, an entirely software implementation or an implementation combining software and hardware aspects. The following detailed description is, therefore, not to be taken in a limiting sense.

**[0028]** Existing search systems retrieve electronic objects, or links to electronic objects, as results. For example, a web search system may retrieve URL links to web pages in response to a query while a file system search interface may retrieve one or more files. Generally, the results returned from existing search systems are not the “precise information” that the user who submitted the query is looking for. For example, in the case of a web search system, the user must click one of the returned URL links to access an electronic object which actually contains the information that the user is looking for.

**[0029]** In this disclosure, the term “passage” is used. A passage is a piece of information or content that provides more direct fulfillment of a user’s search intent. For example, in response to Internet search query, rather than just returning a list of URLs, the content of a group of URLs is searched to identify and extract the relevant information to answer the query. A passage, either generated or extracted from the URLs, will be returned to the user. It will be appreciated that, while this disclosure often refers to URLs and returning text based passages, this disclosure is not so limited. A passage is not limited to text information. In various aspects a passage can be an image, video, and/or a combination of different information.

**[0030]** Additionally, this disclosure discusses identifying the semantic meaning of words, phrases, (e.g., any n-gram). As a specific example, semantic meaning may be interpreted using keywords and entities between a query and a passage. Keywords are words that may be helpful in determining intents and domains of sentences, phrases, and passages in an electronic object as well as a query. Entities are similar to keywords, but additionally are words or phrases that have, as used in the context of the natural language query or passage, an alternative meaning than the literal definition. For example, the words “super bowl” typically would not literally mean an excellent stadium, but normally refers to the championship game of the National Football League.

**[0031]** Aspects of the present technology include the use of machine-learned models. The machine learned models will use a set of training data to learn how to recognize semantic patterns in natural language. The training data will be tagged in such a way as to train the model for a particular task. For example, when identifying key words in a sentence, the model will be trained to recognize California king as one entity related to mattresses, and not necessarily a monarch of the state of California. One of skill in the art will appreciate that various different types of machine learning approaches and models can

be employed by the technology disclosed herein including, but not limited to, decision tree learning, association rule learning, artificial neural networks, deep learning, support vector machines, Bayesian networks, etc.

**[0032]** As a non-limiting example, a user may submit the following query to an Internet search engine: “Will a California king fit on a king frame?” In response to receiving this query, a traditional search engine will provide a list of links to web pages that were identified as related to the query. Rather than, or in addition to, providing the most relevant links based on a key-word search, aspects of the technology described herein returns an answer such as “A standard king-sized bed frame measures 76 inches wide by 80 inches long. A California king measures 72 inches wide and 85 inches long, making it five inches longer but four inches narrower than a standard king-size mattress.” Thus, the technology described herein may be applied to identify/generate a passage that provides a more narrowly tailored answer to this and other queries, thereby making the precise information available to the user in faster than traditional search systems.

**[0033]** Aspects disclosed herein relate to a retrieval system for finding the relevant passages from across an entire web of electronic objects, such as the Internet. Additionally, electronic object ranking and passage ranking algorithms and techniques are disclosed. As such, in addition to providing other benefits, aspects of the technology disclosed herein not provide improvements to the coverage of existing answer retrieval systems but also provide more precise passages by using web ranking algorithms and techniques. For ease of illustration, various aspects of the disclosure are described with respect to providing passages identified or synthesized using web content. However, one of skill in the art will appreciate that the technologies described herein can be employed across various different types of content stores such as, but not limited to, databases, file systems, email mailboxes and/or archives, social media networks, etc.

**[0034]** Turning now to Figure 1, Figure 1 illustrates a networked-computing environment 100 for retrieving relevant passages from a corpus of electronic objects. The corpus of electronic objects is defined by the electronic objects that are accessible by a search application. The size of the corpus of electronic objects may vary depending on the type of search application. For example, the corpus of documents available to a file system search application is limited to the electronic objects stored on the file system. The corpus of electronic objects available to a web search application, however, may include

every electronic object that is accessible on the Internet. As illustrated, Figure 1 includes a computing device 102, a networked-database 104, and a server 106, each of which is communicatively coupled to each other via a network 108.

**[0035]** The computing device 102 may be any suitable type of computing device. For example, the computing device 102 may be one of a desktop computer, a laptop computer, a tablet, a mobile telephone, a smart phone, a wearable computing device, or the like. Additionally, aspects of the current technology include the computing device storing a browser 110 and a search application 112.

**[0036]** A browser 110 may be an internet search browser. A user may input search queries to the browser 110 through a variety of means including text, touch, gesture, or spoken language. The browser 110 may be configured to receive the query and transmit the query to a search engine stored on one or more servers, such as search engine 114 stored on server 106.

**[0037]** As illustrated, the computing device 102 includes a search application 112. The search application 112 may be a search application programmed to receive queries for searching electronic objects such as emails, video files, word processing documents, etc. In certain aspects, the received input is a natural language query. The received query may be sent to the server 106. Aspects of the technology include the search application 112 having the ability to receive queries through text, touch, and/or speech input. The search engine 114 may be configured to receive the query and transmit the query to a search engine stored on one or more servers, such as search engine 114 stored on server 106.

**[0038]** System 100 may also include a database 104. The database 104 may be used to store a variety of information to aid in retrieving relevant passages from a corpus of electronic objects. For example, a Deep Neural Network 118 (“DNN”) and a machine learned model 120 may be provided to the server 106 to aid in retrieving relevant passages from a corpus of electronic objects.

**[0039]** As illustrated, a search engine 114 and a passage extraction engine 116 may reside on a remote device, such as server 106. In other examples, however, the search engine 114 and a passage extraction engine 116 may be stored on another computing device, such as computing device 102. One or both of the search engine 114 and the passage extraction engine 116 receives queries from computing devices such as computing device 102. The search engine 114 receives a query and, using the query, identifies a

corpus of potential electronic objects that may satisfy the query. In aspects of the technology, the passage extraction engine 116 identifies, analyzes, and ranks the passages within a corpus of electronic objects, such as those objects identified by the search engine 114.

**[0040]** Network 108 facilitates communication between devices, such as computing device 102, database 104, and server 106. The network 108 may include the Internet and/or any other type of local or wide area networks. Communication between devices allows for the exchange of queries, information related to the corpus of electronic objects, relevant passages, and other information.

**[0041]** Figure 2 illustrates a method 200 for providing a passage. In examples, method 200 may be executed on a device comprising at least one processor configured to store and execute operations, programs or instructions. However, method 200 is not limited to such examples. In other examples, method 200 may be performed on an application or service for providing the storage and/or manipulation of user input. In at least one example, method 200 may be executed (e.g., computer-implemented operations) by one or more components of a distributed network, such as a web service/distributed network service (e.g. cloud service). Method 200 begins with receive query operation 202. In receive query operation 202, a query is received. The query may be a natural language query received by a search application, a web browser, a personal digital assistant, etc. In other aspects, the query is a structured query that conforms to a specific syntax.

**[0042]** Method 200 proceeds to identify corpus operation 204. In operation 204, a corpus of electronic objects, such as a database storing emails or an Internet domain is identified. In one aspect, a conventional search engine may be employed to identify the corpus of electronic objects relevant to the query. For example, a list of top URLs generated by a search engine may constitute the corpus of electronic objects, with the electronic objects being the page represented by the URLs. In certain aspects, the corpus may be identified based upon the query. Alternatively, the corpus may be identified means without using a query, such as a specific database, a web domain, an email server, etc. In further aspects, the corpus may be identified based upon the networks and/or data stores that are accessible to the application that received the query. In further aspects, additional processing may be performed to identify the corpus. In such aspects, filtering and/or ranking may be utilized to identify a corpus of electronic objects. For example, a

search engine, such as an Internet search engine, may receive a query and, based upon the query, produces a ranked list of top candidate URLs that include information to likely satisfy the query. In such an example, the top candidate URLs, as opposed to all related URLs, may be identified as the corpus of electronic objects. The query used by the search engine may be the same or similar to the query received in operation 202. That is, the search engine process or otherwise alter the query prior to identifying the corpus of electronic objects.

**[0043]** Method 200 then proceeds to parse content operation 206. In one aspect, each electronic object that is part of the corpus may be parsed. In alternate aspects, only a number of highly ranked electronic objects (e.g., 5, 10, 15, 20 . . . 200, etc.) are chosen from the ranked corpus identified in operation 204. Because parsing may be computationally expensive and time consuming, it may be beneficial to only parse a set of highly ranked electronic objects in order provide a timely response to the query. Alternatively, pre-parsing of the corpus of electronic objects may be performed before receiving the query. At operation 206, the content of the electronic objects are parsed. For example, where the corpus is web pages that have been ranked ordered by a search engine, the content of each web page is parsed. Parsing may determine and delineate the content into passages. This may be a rule based system, such as delineating content into passages by identifying carriage returns in a document, the presence of a body, chapters in a video, a digital mark in an audio clip, etc. Additionally, or alternatively, content with similar semantic meaning may be grouped together as a passage.

**[0044]** Additionally, a semantic meaning of each passage may be determined during the parsing operation 206. The semantic meaning may be determined using a machine learned model.

**[0045]** Method 200 then proceeds to rank passages in each electronic object (or a subset of electronic objects identified by rank) in the corpus at operation 208. Ranking the passages of each electronic object may be performed by comparing the semantic meaning of the passage with the semantic intent of the query. Key entities of a query may be compared to entities in the passage by building a semantic translation model between the passage and the query. Alternatively, or additionally, the query answer type may be compared to the information presented in the passage. Still further, the context of the query and the passage may be compared. Further, key-features between the query and the

passage may be analyzed. The result of performing one or more of these techniques on a passage is a numerical score indicating the likelihood that the passage directly answers the received query.

**[0046]** Method 200 then proceeds to rank aggregate passages operation 210. The method 200 may identify each ranked passage from each top candidate document. Each ranked passage may then be aggregated with the ranked passages from other electronic objects. In aspects, the aggregated passages are be ranked against one another. In aspects, ranking occurs by comparing the numerical scores of the passages, which score was determined in operation 208. Alternatively, an alternate model may be used to generate a new score for the aggregated passages.

**[0047]** Method 212 then proceeds to select passage operation 212. In select passage operation, the one or more top ranked passages from the ranked aggregated passages is chosen. The selected passage(s) is then returned at operation 214.

**[0048]** Figure 3 illustrates an exemplary passage retrieval system 300 for returning relevant passages based on a query. The passage retrieval system 300 may be integrated with a web search infrastructure. Passage retrieval system 300 interprets a query at interpretation service 302. In query interpretation service 302, the query may be interpreted to determine the intent of the query as well as an expected answer type. For example, user may query a system with “what drinks have the least amount of calories.” The intent of the user may be identified as wanting to search the Internet to determine a list of beverages along with the number of calories associated with each beverage, rank ordered from least to highest. The intent of the query may be determined using machine learning, a deep neural network, or any other mechanism known to the art without departing form the scope of this disclosure. Query interpretation service 302 may also perform additional operations such as parsing the query, transforming the query, determine whether additional information is needed for the query, prompt the user for additional information, etc. The query, along with any interpretations, are sent to selection service 304.

**[0049]** At passage extraction service 303, one or more electronic objects are analyzed to determine for potential relevant passages. The extraction service may analyze a corpus of electronic documents prior to any query being received. The extraction service may analyze electronic objects by identifying signals within the object. The signals suggest

whether the electronic object has content that may be useful. In embodiments, the content is parsed into individual passages. Where the passage has a signal that associated with likely relevance, the passage is flagged as potentially relevant. In aspects, the passage extraction service identifies and formulates passage candidates from electronic objects that may be relevant (or will not be relevant) prior to receiving a query (offline extraction) or after receiving a query (online extraction). The offline extraction allows for faster processing of a query, in aspects. As an example, a web domain may be searched and the content may be parsed. The content may be identified as containing mostly advertisements, or empty pages, in such an example, the web pages may be marked as not relevant. The flagged passages and/or electronic objects are sent to passage ranking service 310.

**[0050]** At selection service 304, candidate electronic objects are determined for further processing. Continuing with the example above, the system may identify a list of web pages that are sorted by relevant keyword calorie, drink, beverage, etc. Other methods of identifying relevant webpages in the URL may be used. The resulting corpus of electronic documents is sent to the document ranking service 306. The returned electronic objects may then be ranked by keyword or semantic relevance at document ranking service 306.

**[0051]** The top-most candidates determined at document ranking service 306 are passed to pre-caption ranking service 308. Pre-caption ranking service 308 further determines a relevance of the electronic objects. A numerical score, such as a probability, may be associated with each candidate document in order to determine relevance. The result of pre-caption ranking service 308 is a list of pre-caption ranked electronic objects (such as URLs linking relevant web pages). The results of pre-caption ranking service 308 are passed to both caption generator 311 and passage ranking service 310.

**[0052]** At passage ranking service 310, individual passages of each of the results of pre-caption ranking service 308 are ranked. The identification of passages is performed with passage extraction service 303. To continue with the above example, the pre-caption ranking service may have returned a list of 10 URLs that are relevant to the intent of the user. Each passage within the electronic objects identified by each URL is ranked. The different electronic objects identified by the different URLs may contain a varying number of passages. That is, the electronic object identified by URL 1 (hereinafter referred to as "URL 1") may have 3 passages, the electronic object identified by URL 2 (hereinafter

referred to as “URL 2”) may have 7 passages, etc. The passages within each identified electronic object may be ranked against one another. As an example, the URL 1 will be ranked against the other passages of URL 1, the passages in URL 2 will be ranked against other passages in URL 2, and so on. Ranking may be performed using keywords, relevance to the user intent, etc.

**[0053]** The passage ranking service 310 identifies and ranks passages from the electronic objects included in the pre-caption results. Ranking the passages of each page may be performed by comparing the semantic meaning of the passage with the semantic intent of query. Key entities of a query may be compared to entities in the passage by building a semantic translation model between the passage and the query. Also, the query answer type may be compared to the information presented in the passage. Further, the context of the query and the passage may be compared. Additionally, cross-document aggregation may be used. Further, key-features between the query and the passage may be analyzed. In aspects of the technology, the result of performing one or more of these techniques on a passage is a numerical score indicating the likelihood that the passage directly answers the received query.

**[0054]** The n-top ranked passages of each electronic object ranked at pre-caption ranking service 310 is passed to aggregate rank service 312. At aggregate rank service 312, all of the ranked passages are ranked against each other. For example, each of the top three passages from URL 1 may be received, the top three passages from URL 2 may be received, and the top three passages from URL 3 may be received for a total of 9 URLs. These top passages, known as the aggregated passages, are then ranked against each other. Ranking may occur using the numerical scores assigned to each passage at service. Additionally, or alternatively, the numerical scores of each passage may be recalculated. The resulting top n-most passage is passed to merge service 312.

**[0055]** Returning to caption generator 311, caption generator generates one caption for each of the pre-caption ranked electronic objects resulting from the pre-caption ranking service 308. The caption generator parses the content of each pre-caption ranked electronic object (or portion thereof) and generates a caption. The caption may be generated using a machine learned model. The caption is applied the electronic objects and to form a captioned electronic objects. The captioned electronic objects are then sent to the post-caption ranking service 314.

**[0056]** The post-caption re-ranking service 314 re-ranks the ranked electronic objects received by the caption generator 311. The re-ranking includes an analysis of the summary generated in caption generator 311. The analysis may be performed by a machine learned model. This results in post-caption re-ranked electronic objects. The results are sent to merge service 316.

**[0057]** In aspects, at merge service 316, the top n-most passage(s) are formatted for presentation along with one or more links to post-caption re-ranked electronic objects. For example, the top most passage from an electronic object (such as a web page) may be formatted in such a way so as to display in a mobile device, and a link to the top post-caption re-ranked electronic objects may be included.

**[0058]** Figure 4 illustrates an example output of a relevant passage and a link to an exemplar electronic object. As illustrated example output 400 includes a passage 402 and a link 404.

**[0059]** The passage 402 is a passage that answers a query, such as a user entering a query into a browser, based on the interpretation of the intent of the query. The summary may be an actual passage from one or more top electronic objects, such as the web page linked by link 404, or may be generated using an amalgamation of various top-ranked electronic objects.

**[0060]** The link 404 shown in example output 400 is a URL, though it need not be. In some embodiments, multiple links are provided. The multiple links may be the top-n most relevant links as determined by the systems and methods described herein.

**[0061]** Figure 5 illustrates a method 500 for identifying a query dependent passage. In aspects, method 500 may be implemented to find relevant passages based on comparing a passage to the query. In examples, method 500 may be executed on a device comprising at least one processor configured to store and execute operations, programs or instructions. However, method 500 is not limited to such examples. In other examples, method 500 may be performed on an application or service for providing the storage and/or manipulation of user input. In at least one example, method 500 may be executed (e.g., computer-implemented operations) by one or more components of a distributed network, such as a web service/distributed network service (e.g. cloud service).

**[0062]** Method 500 begins with receive query operation 502. In receive query operation 502, a query is received. The query may be received from a client device, such as a client device 102. The query may be a natural language query entered into a web browser. The query is then parsed to determine information about the query. For example, the semantic meaning of the query may be determined, keywords and entities may be identified, and answer type may be determined. This may be performed using a machine learned model.

**[0063]** Method 500 then proceeds to receive electronic object operation 504. In operation 504, a list of electronic objects (such as web pages, emails, word documents) is received based on the received query. For example, where the query is an Internet search query entered into a browser, a search engine will run a search using the query. In such a case, a list of links will be received.

**[0064]** Method 500 then proceeds to parse electronic object operation 506. In operation, the content of the electronic object is parsed. The parsing may be done by extracting the content (textual, audio, video, etc.) of the electronic object, such as a web page, email, or word processing documents. The parsing may be done using a machine learned model. Parsing may identify passages in the content, such as paragraphs or sentences in an electronic object that have similar semantic meaning. From such parsing, the semantic meaning of the passages within the electronic object, is identified.

**[0065]** The semantic meaning and other information derived at operation 506 may be compared to information related to the query at the compare operation 508. For example, the keywords, entities, answer type derived from in operation 502 may be compared to the information derived at operation 506. The comparison may identify passages that are highly relevant to the query. Each passage may be ranked according to relevance of the passage.

**[0066]** The method 500 then proceeds to identify passage based on comparison operation 510. In aspects, the top n-most passages are identified as highly relevant. Relevancy may be determined in a variety of ways. For example, key entities of a query may be compared to entities in the passage by building a semantic translation model between the passage and the query. Also, the query answer type may be compared to the information presented in the passage. Further, the context of the query and the passage may be compared. Additionally, cross-document aggregation may be used. Further, key-features between the query and the passage may be analyzed. In aspects of the technology,

the result of performing one or more of these techniques on a passage is a numerical score indicating the likelihood that the passage directly answers the received query.

**[0067]** One or more highly relevant passages may be identified in identify operation 510. Identifying may be performed by choosing the top most relevant passage. The passages may have been assigned a numerical score. In such cases, identification may be any passage that exceeds a threshold numerical value.

**[0068]** Figure 6 illustrates a method 600 for identifying a query independent passage in a document. In aspects, method 600 may be used to identify passages as potentially relevant irrespective of any query. That is, no query is necessarily used to determine potential relevance. In examples, method 600 may be executed on a device comprising at least one processor configured to store and execute operations, programs or instructions. However, method 600 is not limited to such examples. In other examples, method 600 may be performed on an application or service for providing the storage and/or manipulation of user input. In at least one example, method 600 may be executed (e.g., computer-implemented operations) by one or more components of a distributed network, such as a web service/distributed network service (e.g. cloud service).

**[0069]** Method 600 begins with identify electronic objects operation 602. At operation 602, various electronic objects may be identified. Identification may be based on a set list (such as all web pages in a popular URL, all emails from a particular user, all documents created on or before a certain date).

**[0070]** Method 600 then proceeds with parse content operation 604. In parse content operation, the electronic objects are parsed to determine the attributes of the electronic object. Parsing may include identifying metadata for the corpus of electronic objects. In other aspects, the content may be parsed. For example, where the electronic object is a web page, it may be determined that the URL does not include any discernable information, is mostly dedicated to advertisements, or is associated with a domain.

**[0071]** Method 600 then proceeds to compare parsed content to static signal operation 606. Static signals may include a variety of signals. For example, where the electronic object is a document, the static signal may be a threshold length of the document, a threshold number of grammatical errors in a document, a list of trusted/untrusted authors, etc. Where the electronic object is a URL, the static signal may be untrusted/trusted URLs, the number of pop-ups, the presence of a paywall, etc. Electronic objects or passages

within the electronic object may be flagged or ranked as useful/not useful based on the presence or absence of a static signal.

**[0072]** Method 600 then proceeds to generate or identify passage 608. In aspects, passages of the electronic object that are flagged as not useful may be excluded. Where the static signal highlights highly relevant passages, those passages may be identified as useful and flagged as much.

**[0073]** Figure 7 illustrates a method 700 for ranking a passage based on a semantic translation model. In examples, method 700 may be executed on a device comprising at least one processor configured to store and execute operations, programs or instructions. However, method 700 is not limited to such examples. In other examples, method 700 may be performed on an application or service for providing the storage and/or manipulation of user input. In at least one example, method 700 may be executed (e.g., computer-implemented operations) by one or more components of a distributed network, such as a web service/distributed network service (e.g. cloud service).

**[0074]** Method 700 begins with receive query 702. In receive query operation 702, a query is received. The query may be received from a client device and sent to a server for processing, such as a client device 102. The query may be a natural language query entered into a web browser.

**[0075]** Method 700 then proceeds to identify semantic units in a query operation 704. Semantic units in the query operation may be identified using a machine learned model. For example, for the query “Will a California king fit on a king frame?” the entities in query may be identified as {California\_king} and {king\_frame}.

**[0076]** The method 700 then proceeds to identify semantic units in a passage 706. The semantic units of a passage may be identified using a machine learned model. Following the above example, a passage may be: “Insta-lock Queen, King, Cal-king Bedframe It's a bed frame We purchased this bed frame for a mattress & box springs. The instructions were easy to read & the frame went together without any problems.” Another passage may be “A California king bed measures 72 inches wide and 84 inches long, and a standard king mattress measures 76 inches wide and 80 inches long, so while a California king mattress is four inches longer, it is also about four inches thinner than a standard king bed.” Each passage is parsed to identify entities.

**[0077]** Method 708 then proceeds to generate semantic translation model “TM” 708. In operation 708 a TM model is generated by building a semantic unit relationship table with translation probabilities between the query and the passage. The entity relationship table, in aspects, generates a probability that the entity identified in the query is present in the passage.

**[0078]** Figure 8 illustrates experimental output resulting from identifying features of a passage based on a query and a sample target electronic object. The semantic translation model is generated with machine learning method. As illustrated, column one 802 includes semantic units of source query, column two 804 includes target passage semantic units, column three 806, is the probability that the semantic unit in the passage is the same as the entity used in the query, and column 4 808 is the probability that the semantic unit in the query is the same as the semantic unit in the passage. The higher the probability the most likely the source will match the target. Note that  $P(\text{target}|\text{source})$  normally is not the same as  $P(\text{source}|\text{target})$ .

**[0079]** As illustrated, column one 802 includes a semantic unit of source query, column two 804 includes target passage semantic unit, column three 806, is the probability that the semantic unit in the passage is the same as the semantic unit used in the query, and column 4 808 is the probability that the semantic unit in the query is the same as the semantic unit in the passage.

**[0080]** Figure 9 is a method 900 of matching a passage type to a query type. In examples, method 900 may be executed on a device comprising at least one processor configured to store and execute operations, programs or instructions. However, method 900 is not limited to such examples. In other examples, method 900 may be performed on an application or service for providing the storage and/or manipulation of user input. In at least one example, method 900 may be executed (e.g., computer-implemented operations) by one or more components of a distributed network, such as a web service/distributed network service (e.g. cloud service).

**[0081]** Method 900 beings with extract answer type from query operation 902. In aspects, a query type of the answer may be numerical, a location, a time, a commercial center, etc. For example, a query may ask “how many teaspoons in a cup.” In such an instance, the answer to a query type would be numerical. Additionally, a query may be “where can I buy sushi grade fish near me.” In such a case, the answer type would be a

commercial center, such as a fish store. In aspects, a query is parsed using a machine learned model to determine the answer-type of the query.

**[0082]** The method 900 then proceeds to extract entity type from passage 904. An entity type may be identified an extracted using a machine learned model. The entity type may be one that is numerical, a location, a time, a commercial center, etc. For example, if the passage includes the phrase “There are 48 teaspoons in a cup,” the entity 48 would be identified as numerical. Another passage may include the phrase “teaspoon sizes originated in England, where the price of tea caused tea cups and spoons to shrink.” The entity England would be identified as a location. An entity may have multiple entity types. For example, consider the following text from a passage: “Von Miller was the MVP of Super Bowl 50.” The following types may be associated with Von Miller: Person, Athlete, American Football Player.

**[0083]** The method 900 then proceeds to rank passages 906. In rank passage operation 906, the passages are ranked by comparing the answer type to the entity type identified in operation 904. Following the example above, if the query is expected to return a numerical answer (e.g., “how many teaspoons in a cup”) passages with a numerical entity (e.g., “there are 48 teaspoons in a cup”) would be flagged as relevant. Other passages where the entity did not match the answer type (e.g., “teaspoon sizes originated in England, where the price of tea caused tea cups and spoons to shrink”) would be flagged as likely not relevant. In aspects, the ranking of the passages is based on a machine learned model.

**[0084]** Figure 10 is an example output using a method of matching a passage type to a query type. As illustrated, a query 1002 “what are the drinks with the` least amount of calories” has been parsed to identify a semantic intent 1004 “alcoholic drinks with least calories.” Fig 10 also includes the resolution of the answer type 1006 sought by the query 1002 as numeric and has identified key words as “calories.” Identification of semantic intent, answer types, and keywords may be performed by using a machine learned model.

**[0085]** Figure 10 also includes a passage 1008 that has been parsed. In aspects, the passage is parsed using a machine learned model. As illustrated, the parsed passage includes identified entities “calories” that is the same as the key words identified in the query. Further, the calories include a modifier “100,” which matches up with the query answer type numeric. Accordingly, the passage will be given a relatively high score. In aspects, the score is calculated by adding the number of like elements of a passage (e.g.,

like answer types/entities, and keywords) of a passage. Scores may be calculated other ways.

**[0086]** Figure 11 is a method 1100 of matching an electronic object to a query based on contextual meaning. In examples, method 1100 may be executed on a device comprising at least one processor configured to store and execute operations, programs or instructions. However, method 1100 is not limited to such examples. In other examples, method 1100 may be performed on an application or service for providing the storage and/or manipulation of user input. In at least one example, method 1100 may be executed (e.g., computer-implemented operations) by one or more components of a distributed network, such as a web service/distributed network service (e.g. cloud service).

**[0087]** Method 1100 begins with extract nearest heading to passage as at operation 1102. It will be appreciated that heading may refer to any marker in an electronic object that that indicates a section of the electronic object. This may include title, caption, legend, subtitle, subheading, rubric, headline, etc. In operation 1102, the nearest heading to a passage may be extracted. This may be done by identifying tagged content in an electronic object (such as a web page that identifies a phrase as a heading), the title of a video, the subject line of an email, etc. Identification may be performed using a machine learned model.

**[0088]** Method 1100 then proceeds to identify semantic meaning of query 1104. The semantic meaning of a query may be identified using a machine learned model.

**[0089]** Method 1100 then proceeds to compare heading to query operation 1106. In operation 1104, the contextual meaning of each heading is compared to the semantic meaning of the query. Where the heading is similar to the query, the heading is marked as relevant.

**[0090]** Method 1100 then proceeds to flag passage operation 1108. At operation 1104, passages associated with a heading that has been marked as relevant are flagged as likely relevant. In some embodiments, a numerical score associated with the relevance of the passage is increased based on a flagged header.

**[0091]** Figure 12 is an example of matching an electronic object to a query based on the heading of the electronic object. As illustrated, the query 1202 includes “how to change the excel cell width.” The headings of a web page (an example electronic object) are

analyzed to determine potentially relevant passages. The heading and query may be analyzed using the methods described herein, including method 1100. In this example, the title of the heading of the web page 1204 is “Changing Column Widths and Row Heights in Excel.” Further sub-headings within the web page are analyzed. In this case, the sub heading “how to change column widths” 1206 and sub-heading “how to change row heights” 1208 are analyzed. The passage 1208 associated with the sub-heading “how to change column widths” 1206 is flagged as potentially relevant. This relevant passage 1210 may then be processed further, given a numerical score, or have the numerical score increased.

**[0092]** Figure 13 is a method 1300 of ranking passages. Method 1302 begins with select top passages. In operation 1300, top candidate passages are identified.

**[0093]** The method 1300 then proceeds select top terms operation 1304. From each passage 1302. In operation 1304, the terms that are the same or similar to terms of query are identified in each passage. Identification of top terms may be done by using a machine learned model.

**[0094]** The method 1300 then proceeds to identify entities operation 1306. In operation 1306 entities that are similar to or the same as those of the query (or related to those of the query) of the passage are identified. Same, similar, or related entities may be determined using a machine learned model.

**[0095]** The operation then proceeds to rank operation 1308. In operation 1308, each passage is ranked. In an aspect, the passage is ranked summing the number of keywords and entities identified in operation 1306 and 1304. The passage with the highest total may be ranked first, in aspects.

**[0096]** Figure 14 is an example 1400 of applying the method 1300 on the passage “The Fifteenth Amendment (Amendment XV) to the United States Constitution prohibits the federal and state governments from denying a citizen the right to vote based on that citizen's “race, color, or previous condition of servitude” 1402. As illustrated, the query 1404 is “15<sup>th</sup> amendment definition.” In top terms identification 1406, the top terms, such as the term “amendment”, “constitution”, “citizens”, “fifteenth”, and “prohibits” are identified. This is five terms.

**[0097]** The example 1400 also includes identifying similar, same, or related top entities. In the example, the entities “condition\_of\_servitude,” “previous\_condition,” “the\_15th\_amendment to\_the\_united\_states the\_united\_states\_constitution,” “the\_u\_s\_constitution voting\_rights an\_amendment the\_fifteenth\_amendment” “the\_right\_to\_vote southern\_states the\_13<sup>th</sup>” are identified. This is a total of 5 entities. In embodiments, the passage would be given a rank of 10.

**[0098]** Figure 15 illustrates a method 1500 of matching answer patterns in passages to answer types related to queries. Method 1500 begins with determine query answer pattern operation 1502. A query answer pattern is determined based on the query pattern. A query may be in a pattern such as who is noun, where is noun, etc. Using a machine learned model, a translation table between the query pattern and the pattern of the expected answer, e.g., the query answer pattern, may be established.

**[0099]** Method 1500 then proceeds to determine passage pattern operation 1504. In determine passage pattern operation 1504, the pattern of passage is determined. The semantic pattern of the passage may be determined using a machine learned model.

**[00100]** Method 1500 proceeds to compare passage pattern with query answer pattern operation 1506. In this example, the patterns of the passage are identified that have a similar or same pattern as the query answer pattern.

**[00101]** Method 1500 proceeds to score passage operation 1508. Based on the similarity, the passage is scored and/or moved to a higher rank. Various different probabilistic models may be employed to score the passage based upon the similarity.

**[00102]** Figure 16 is an example 1600 of the results of performing method 1500 on a passage.

**[00103]** As illustrated, the query 1602 the query is “who was the king of England after Queen Elizabeth I.” In this example, the query answer pattern is passages with answer pattern 1608 “<answer> was king of *king* <answer>, *who father*, <answer>.”

**[00104]** Top answer 1610 is discovered by analyzing the top N passages 1606. In this example, the answer includes the entities “the\_virgin\_queen anne\_boleyn the\_daughter king\_henry\_viii,” “when\_she queen\_elizabeth\_i prince\_james,” “james\_vi\_of\_scotland,” and “44\_years\_of rule.”

**[00105]** The most relevant passage 1612 associated with the answer is “After the death of Queen Elizabeth I without issue, in 1603, the crowns of England and Scotland were joined in personal union under King James VI of Scotland, who became James I of England.” This may be provided to a user as a direct answer to the query.

**[00106]** Figs. 17-19 and the associated descriptions provide a discussion of a variety of operating environments in which examples are practiced. However, the devices and systems illustrated and discussed with respect to 17-19 are for purposes of example and illustration and are not limiting of a vast number of computing device configurations that are utilized for practicing aspects, described herein.

**[00107]** Figure 17 is a block diagram illustrating physical components (i.e., hardware) of a computing device 1700 with which examples of the present disclosure may be practiced. In a basic configuration, the computing device 1700 includes at least one processing unit 1702 and a system memory 1704. According to an aspect, depending on the configuration and type of computing device, the system memory 1704 comprises, but is not limited to, volatile storage (e.g., random access memory), non-volatile storage (e.g., read-only memory), flash memory, or any combination of such memories. According to an aspect, the system memory 1704 includes an operating system 1705 and one or more program modules 1706 suitable for running software applications 1750. According to an aspect, the system memory 1704 includes the system 100. The operating system 1705, for example, is suitable for controlling the operation of the computing device 1700. Furthermore, aspects are practiced in conjunction with a graphics library, other operating systems, or any other application program, and are not limited to any particular application or system. This basic configuration is illustrated in Figure 17 by those components within a dashed line 1708. According to an aspect, the computing device 1700 has additional features or functionality. For example, according to an aspect, the computing device 1700 includes additional data storage devices (removable and/or non-removable) such as, for example, magnetic disks, optical disks, or tape. Such additional storage is illustrated in Figure 17 by a removable storage device 1709 and a non-removable storage device 1710.

**[00108]** As stated above, according to an aspect, a number of program modules and data files are stored in the system memory 1704. While executing on the processing unit 1702, the program engines 1706 (e.g., the engines of system 1700) perform processes including, but not limited to, one or more of the stages of the methods used for aggregating and

modeling illustrated in Figures 2,3, 5, 6, 7, 9, 11, 13, and 15. According to an aspect, other program modules are used in accordance with examples and include applications such as electronic mail and contacts applications, word processing applications, spreadsheet applications, database applications, slide presentation applications, drawing or computer-aided application programs, etc.

**[00109]** According to an aspect, aspects are practiced in an electrical circuit comprising discrete electronic elements, packaged or integrated electronic chips containing logic gates, a circuit utilizing a microprocessor, or on a single chip containing electronic elements or microprocessors. For example, aspects are practiced via a system-on-a-chip (SOC) where each or many of the components illustrated in Figure 17 are integrated onto a single integrated circuit. According to an aspect, such an SOC device includes one or more processing units, graphics units, communications units, system virtualization units and various application functionality all of which are integrated (or “burned”) onto the chip substrate as a single integrated circuit. When operating via an SOC, the functionality, described herein, is operated via application-specific logic integrated with other components of the computing device 1700 on the single integrated circuit (chip). According to an aspect, aspects of the present disclosure are practiced using other technologies capable of performing logical operations such as, for example, AND, OR, and NOT, including but not limited to mechanical, optical, fluidic, and quantum technologies. In addition, aspects are practiced within a general purpose computer or in any other circuits or systems.

**[00110]** According to an aspect, the computing device 1700 has one or more input device(s) 1712 such as a keyboard, a mouse, a pen, a sound input device, a touch input device, etc. The output device(s) 1714 such as a display, speakers, a printer, etc. are also included according to an aspect. The aforementioned devices are examples and others may be used. According to an aspect, the computing device 1700 includes one or more communication connections 1716 allowing communications with other computing devices 1718. Examples of suitable communication connections 1716 include, but are not limited to, radio frequency (RF) transmitter, receiver, and/or transceiver circuitry; universal serial bus (USB), parallel, and/or serial ports.

**[00111]** The term computer readable media, as used herein, includes computer storage media. Computer storage media include volatile and nonvolatile, removable and non-

removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, or program modules. The system memory 1704, the removable storage device 1709, and the non-removable storage device 1710 are all computer storage media examples (i.e., memory storage.) According to an aspect, computer storage media include RAM, ROM, electrically erasable programmable read-only memory (EEPROM), flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other article of manufacture which can be used to store information and which can be accessed by the computing device 1700. According to an aspect, any such computer storage media is part of the computing device 1700. Computer storage media do not include a carrier wave or other propagated data signal.

**[00112]** According to an aspect, communication media are embodied by computer readable instructions, data structures, program modules, or other data in a modulated data signal, such as a carrier wave or other transport mechanism, and include any information delivery media. According to an aspect, the term “modulated data signal” describes a signal that has one or more characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media include wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, radio frequency (RF), infrared, and other wireless media.

**[00113]** Figures 18A and 18B illustrate a mobile computing device 1800, for example, a mobile telephone, a smart phone, a tablet personal computer, a laptop computer, and the like, with which aspects may be practiced. With reference to Figure 18A, an example of a mobile computing device 1800 for implementing the aspects is illustrated. In a basic configuration, the mobile computing device 1800 is a handheld computer having both input elements and output elements. The mobile computing device 1800 typically includes a display 1805 and one or more input buttons 1810 that allow the user to enter information into the mobile computing device 1800. According to an aspect, the display 1805 of the mobile computing device 1800 functions as an input device (e.g., a touch screen display). If included, an optional side input element 1815 allows further user input. According to an aspect, the side input element 1815 is a rotary switch, a button, or any other type of manual input element. In alternative examples, mobile computing device 1800 incorporates more or fewer input elements. For example, the display 1805 may not be a

touch screen in some examples. In alternative examples, the mobile computing device 1800 is a portable phone system, such as a cellular phone. According to an aspect, the mobile computing device 1800 includes an optional keypad 1835. According to an aspect, the optional keypad 1835 is a physical keypad. According to another aspect, the optional keypad 1835 is a “soft” keypad generated on the touch screen display. In various aspects, the output elements include the display 1805 for showing a graphical user interface (GUI), a visual indicator 1820 (e.g., a light emitting diode), and/or an audio transducer 1825 (e.g., a speaker). In some examples, the mobile computing device 1800 incorporates a vibration transducer for providing the user with tactile feedback. In yet another example, the mobile computing device 1800 incorporates input and/or output ports, such as an audio input (e.g., a microphone jack), an audio output (e.g., a headphone jack), and a video output (e.g., a HDMI port) for sending signals to or receiving signals from an external device. In yet another example, the mobile computing device 1800 incorporates peripheral device port 1840, such as an audio input (e.g., a microphone jack), an audio output (e.g., a headphone jack), and a video output (e.g., a HDMI port) for sending signals to or receiving signals from an external device.

**[00114]** Figure 18B is a block diagram illustrating the architecture of one example of a mobile computing device. That is, the mobile computing device 1800 incorporates a system (i.e., an architecture) 1802 to implement some examples. In one example, the system 1802 is implemented as a “smart phone” capable of running one or more applications (e.g., browser, e-mail, calendaring, contact managers, messaging clients, games, and media clients/players). In some examples, the system 1802 is integrated as a computing device, such as an integrated personal digital assistant (PDA) and wireless phone.

**[00115]** According to an aspect, one or more application programs 1850 are loaded into the memory 1862 and run on or in association with the operating system 1864. Examples of the application programs include phone dialer programs, e-mail programs, personal information management (PIM) programs, word processing programs, spreadsheet programs, Internet browser programs, messaging programs, and so forth. According to an aspect, system 100 is loaded into memory 1862. The system 1802 also includes a non-volatile storage area 1868 within the memory 1862. The non-volatile storage area 1868 is used to store persistent information that should not be lost if the system 1802 is powered down. The application programs 1850 may use and store information in the non-volatile

storage area 1868, such as e-mail or other messages used by an e-mail application, and the like. A synchronization application (not shown) also resides on the system 1802 and is programmed to interact with a corresponding synchronization application resident on a host computer to keep the information stored in the non-volatile storage area 1868 synchronized with corresponding information stored at the host computer. As should be appreciated, other applications may be loaded into the memory 1862 and run on the mobile computing device 1800.

**[00116]** According to an aspect, the system 1802 has a power supply 1870, which is implemented as one or more batteries. According to an aspect, the power supply 1870 further includes an external power source, such as an AC adapter or a powered docking cradle that supplements or recharges the batteries.

**[00117]** According to an aspect, the system 1802 includes a radio 1872 that performs the function of transmitting and receiving radio frequency communications. The radio 1872 facilitates wireless connectivity between the system 1802 and the “outside world,” via a communications carrier or service provider. Transmissions to and from the radio 1872 are conducted under control of the operating system 1864. In other words, communications received by the radio 1872 may be disseminated to the application programs 1850 via the operating system 1864, and vice versa.

**[00118]** According to an aspect, the visual indicator 1820 is used to provide visual notifications and/or an audio interface 1874 is used for producing audible notifications via the audio transducer 1825. In the illustrated example, the visual indicator 1820 is a light emitting diode (LED) and the audio transducer 1825 is a speaker. These devices may be directly coupled to the power supply 1870 so that when activated, they remain on for a duration dictated by the notification mechanism even though the processor 1860 and other components might shut down for conserving battery power. The LED may be programmed to remain on indefinitely until the user takes action to indicate the powered-on status of the device. The audio interface 1874 is used to provide audible signals to and receive audible signals from the user. For example, in addition to being coupled to the audio transducer 1825, the audio interface 1874 may also be coupled to a microphone to receive audible input, such as to facilitate a telephone conversation. According to an aspect, the system 1802 further includes a video interface 1876 that enables an operation of an on-board camera 1830 to record still images, video stream, and the like.

**[00119]** According to an aspect, a mobile computing device 1800 implementing the system 1802 has additional features or functionality. For example, the mobile computing device 1800 includes additional data storage devices (removable and/or non-removable) such as, magnetic disks, optical disks, or tape. Such additional storage is illustrated in Figure 18B by the non-volatile storage area 1868.

**[00120]** According to an aspect, data/information generated or captured by the mobile computing device 1800 and stored via the system 1802 are stored locally on the mobile computing device 1800, as described above. According to another aspect, the data are stored on any number of storage media that are accessible by the device via the radio 1872 or via a wired connection between the mobile computing device 1800 and a separate computing device associated with the mobile computing device 1800, for example, a server computer in a distributed computing network, such as the Internet. As should be appreciated such data/information are accessible via the mobile computing device 1800 via the radio 1872 or via a distributed computing network. Similarly, according to an aspect, such data/information are readily transferred between computing devices for storage and use according to well-known data/information transfer and storage means, including electronic mail and collaborative data/information sharing systems.

**[00121]** Figure 19 illustrates one example of the architecture of a system for web-scale passage retrieval as described above. Content developed, interacted with, or edited in association with the system 100 is enabled to be stored in different communication channels or other storage types. For example, various documents may be stored using a directory service 1922, a web portal 1924, a mailbox service 1926, an instant messaging store 1928, or a social networking site 1930. The web-scale passage retrieval system 100 is operative to use any of these types of systems or the like for improving search efficiency, as described herein. According to an aspect, a server 1920 provides the web-scale passage retrieval system 100 to clients 1905a,b,c. As one example, the server 1920 is a web server providing the system 100 over the web. The server 1920 provides the system 100 over the web to clients 1905 through a network 1940. By way of example, the client computing device is implemented and embodied in a personal computer 1905a, a tablet computing device 1905b or a mobile computing device 1905c (e.g., a smart phone), or other computing device. Any of these examples of the client computing device are operable to obtain content from the store 1916.

**[00122]** Implementations, for example, are described above with reference to block diagrams and/or operational illustrations of methods, systems, and computer program products according to aspects. The functions/acts noted in the blocks may occur out of the order as shown in any flowchart. For example, two blocks shown in succession may in fact be executed substantially concurrently or the blocks may sometimes be executed in the reverse order, depending upon the functionality/acts involved.

**[00123]** The description and illustration of one or more examples provided in this application are not intended to limit or restrict the scope as claimed in any way. The aspects, examples, and details provided in this application are considered sufficient to convey possession and enable others to make and use the best mode. Implementations should not be construed as being limited to any aspect, example, or detail provided in this application. Regardless of whether shown and described in combination or separately, the various features (both structural and methodological) are intended to be selectively included or omitted to produce an example with a particular set of features. Having been provided with the description and illustration of the present application, one skilled in the art may envision variations, modifications, and alternate examples falling within the spirit of the broader aspects of the general inventive concept embodied in this application that do not depart from the broader scope.

We claim:

1. A computer implemented method comprising:
  - identifying a corpus of electronic objects for passage extraction;
  - extracting passages from each identified electronic object;
  - receiving a query;
  - ranking extracted passages in each of the electronic objects based on the received query to produce top-ranked passages for each electronic object;
  - aggregating the top-ranked passages of each electronic object to form aggregated passages;
  - ranking the aggregated passages to form ranked aggregated passages;
  - selecting at least one top-ranked passage from the ranked passages;
  - return the at least one top-ranked passage.
  
2. The computer implemented method of claim 1, wherein the electronic objects are one of emails, web pages, images, audio files, videos, or documents.
  
3. The computer implemented method of claim 1, wherein ranking the top-ranked passages comprises:
  - parsing the query to determine query information, wherein the query information comprises at least one of:
    - a query semantic meaning;
    - query keywords; and
    - query entities;
  - parsing the passage to determine passage information, wherein the passage information comprises at least one of:
    - a passage semantic meaning;
    - passage keywords; and

passage entities; and

generating a machine learning passage ranking model with a variety of semantic features, including, a semantic translation model using the query information and the passage information.

4. The computer implemented method of claim 1, further comprising providing a link to the electronic object from where the passage was retrieved.

5. The computer implemented method of claim 4, wherein the passage and the link are provided to a client device.

6. The computer implemented method of claim 4, wherein the query was received from the client device.

7. The computer implemented method of claim 1, wherein the corpus of electronic objects is stored in a distributed network.

8. A computer implemented method comprising:

receiving query;

determining an intent based at least upon the query;

in response to determining the intent, performing a search, using a search engine, to generate search results, wherein the search results include web pages which are ordered according to a rank;

analyzing a plurality of passages from each of at least a subset of the ranked search results to produce a plurality of top-ranked results for each web page;

aggregating the top-ranked results for each web page for the subset of the ranked search results;

ranking the aggregated top-ranked results to identify at least one top-ranked passage;

providing the at least one top-ranked passage.

9. The computer implemented method of claim 8, wherein ranking a passage at least one of the web pages comprises:

analyzing, using a machine learning model, the query to determine the semantic intent of the query;

analyzing, using the machine learning model, the passage;

generating a machine learning passage ranking model with a variety of semantic features, including, a semantic translation model using the query information and the passage information; and

placing the passage in an order among other passages of the plurality of passages based on the numerical score.

10. The computer implemented method of claim 8, wherein the electronic objects are one of emails, web pages, audios, images, videos, or documents.

11. The computer implemented method of claim 8, further comprising providing a link to the web page that contains the at least one top-ranked passage.

12. The computer implemented method of claim 8, wherein the link is returned to a client device.

13. The computer implemented method of claim 12, wherein the query was received from the client device.

14. The computer implemented method of claim 8, wherein the corpus of electronic objects is stored in a distributed network.

15. A computer system comprising:

receiving query;

determining an intent based at least upon the query;

in response to determining the intent, performing a search, using a search engine, to generate search results, wherein the search results include web pages which are ordered according to a rank;

analyzing a plurality of passages from each of at least a subset of the ranked search results to produce a plurality of top-ranked results for each web page;

aggregating the top-ranked results for each web page for the subset of the ranked search results;

ranking the aggregated top-ranked results to identify at least one top-ranked passage;

providing the at least one top-ranked passage.

16. The computer implemented method of claim 15, wherein ranking a passage at least one of the web pages comprises:

analyzing, using a machine learning model, the query to determine the semantic intent of the query;

analyzing, using the machine learning model, the passage;

generating a machine learning passage ranking model with a variety of semantic features, including, a semantic translation model using the query information and the passage information; and

placing the passage in an order among other passages of the plurality of passages based on the numerical score.

17. The computer implemented method of claim 15, wherein the electronic objects are one of emails, web pages, or documents.

18. The computer implemented method of claim 15, further comprising providing a link to the web page that contains the at least one top-ranked passage.

19. The computer implemented method of claim 15, wherein the link is returned to a client device.

20. The computer implemented method of claim 19, wherein the query was received from the client device.




FIG. 1

2/16

200

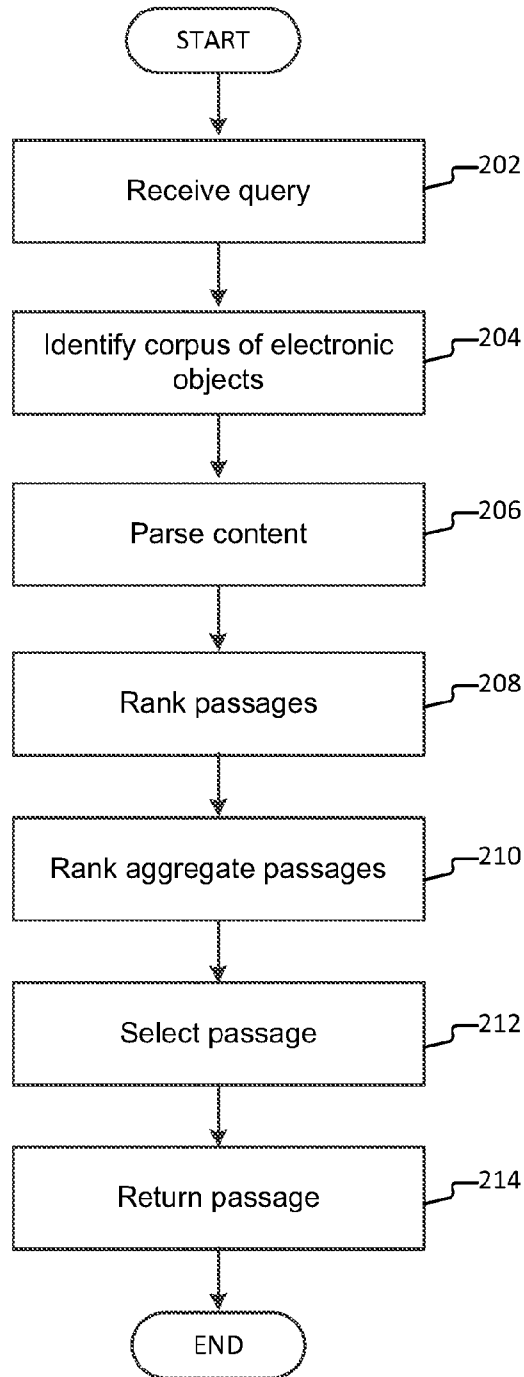


FIG. 2

300

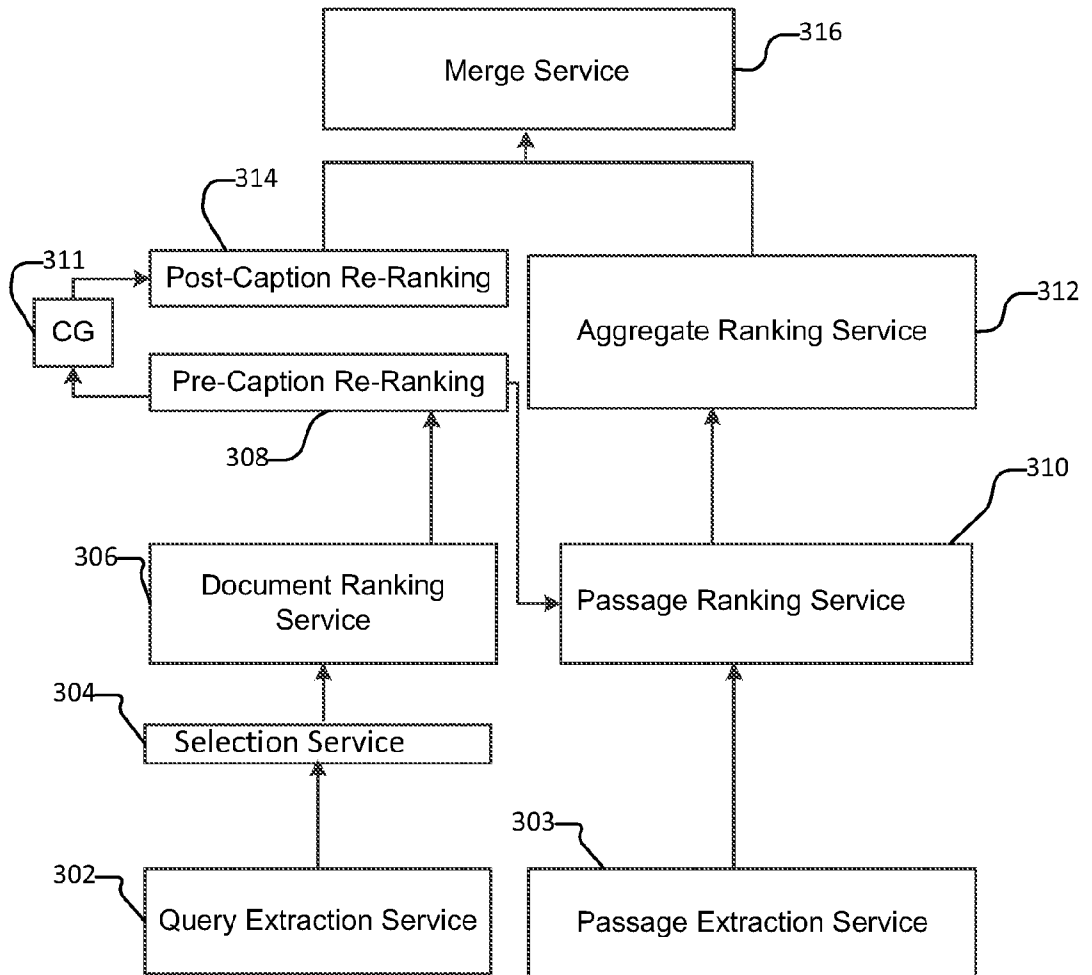


FIG. 3

400




FIG. 4

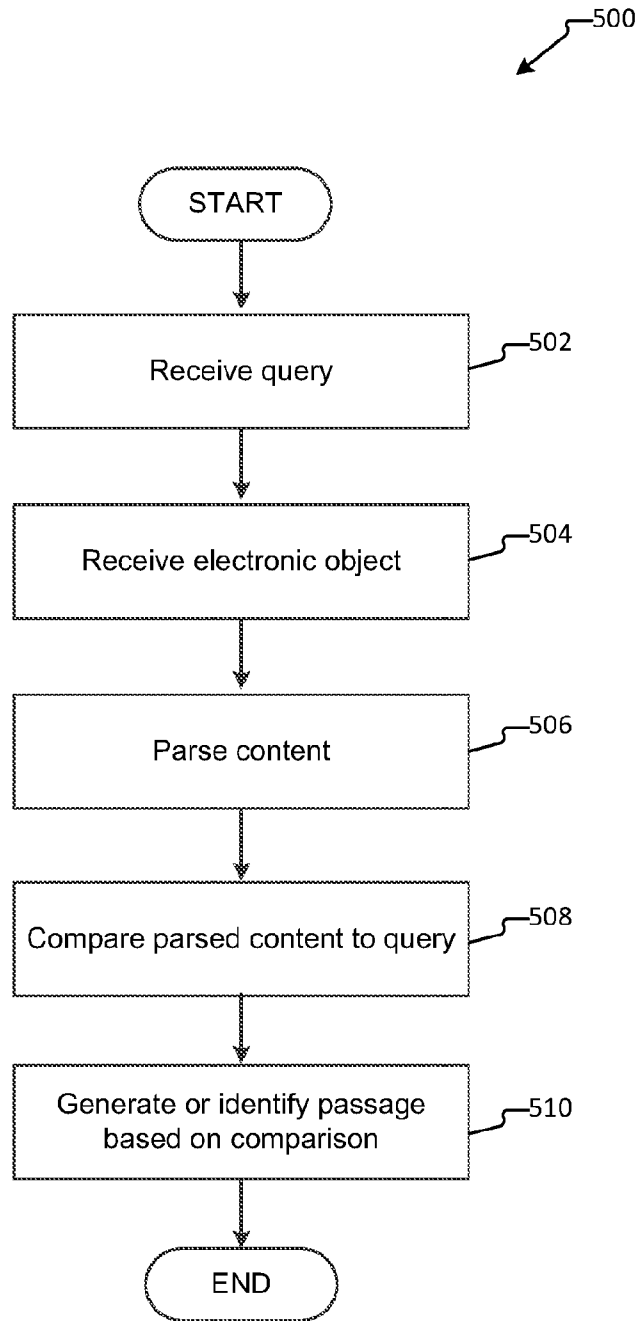


FIG. 5

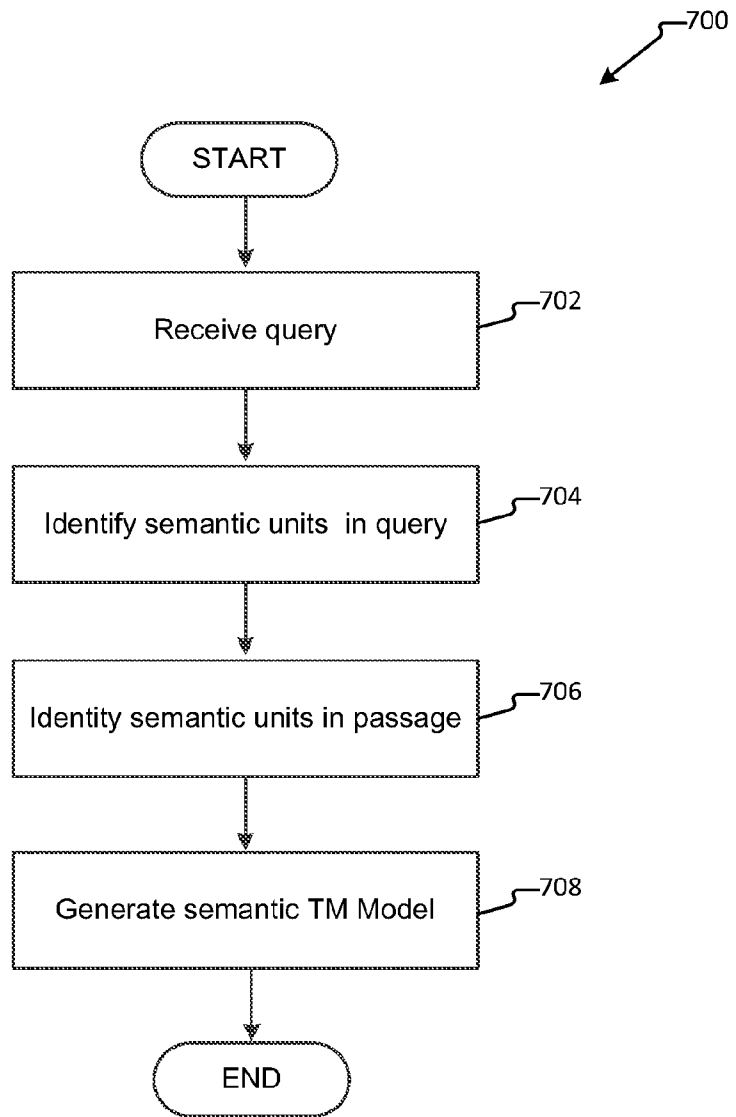


FIG. 7

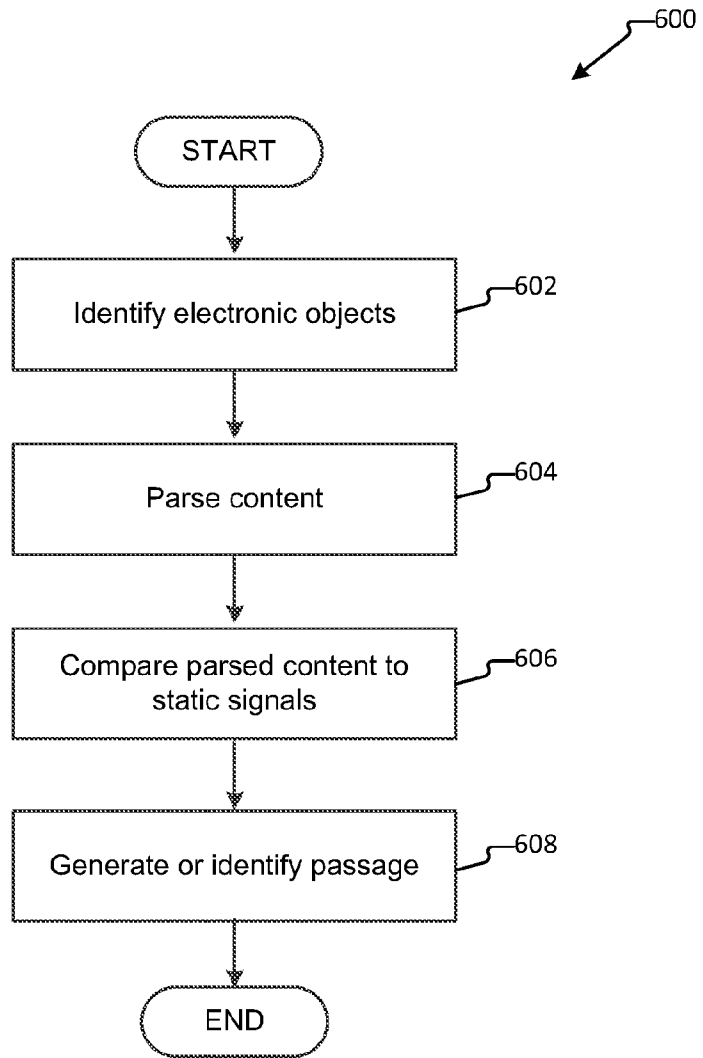


FIG. 6

800

Source (Query)	Target (Passage)	Probability (Target / Source)	Probability (Source / Target)
california_king	california_king_bed	0.00268	0.22764
california_king	cal_king	0.00933	0.11795
california_king	king_mattress	0.00150	0.00314
california_king	long	0.01164	0
california_king	inch	0.00596	0
.....	...	...	...
king_frame	inch	0.22311	0
king_frame	king_mattress	0.12194	0.02997
king_frame	bed_frame	0.12277	0
.....	...	...	...

FIG. 8

Key Feature:  
Query Answer Type

1000





FIG. 9

FIG. 10

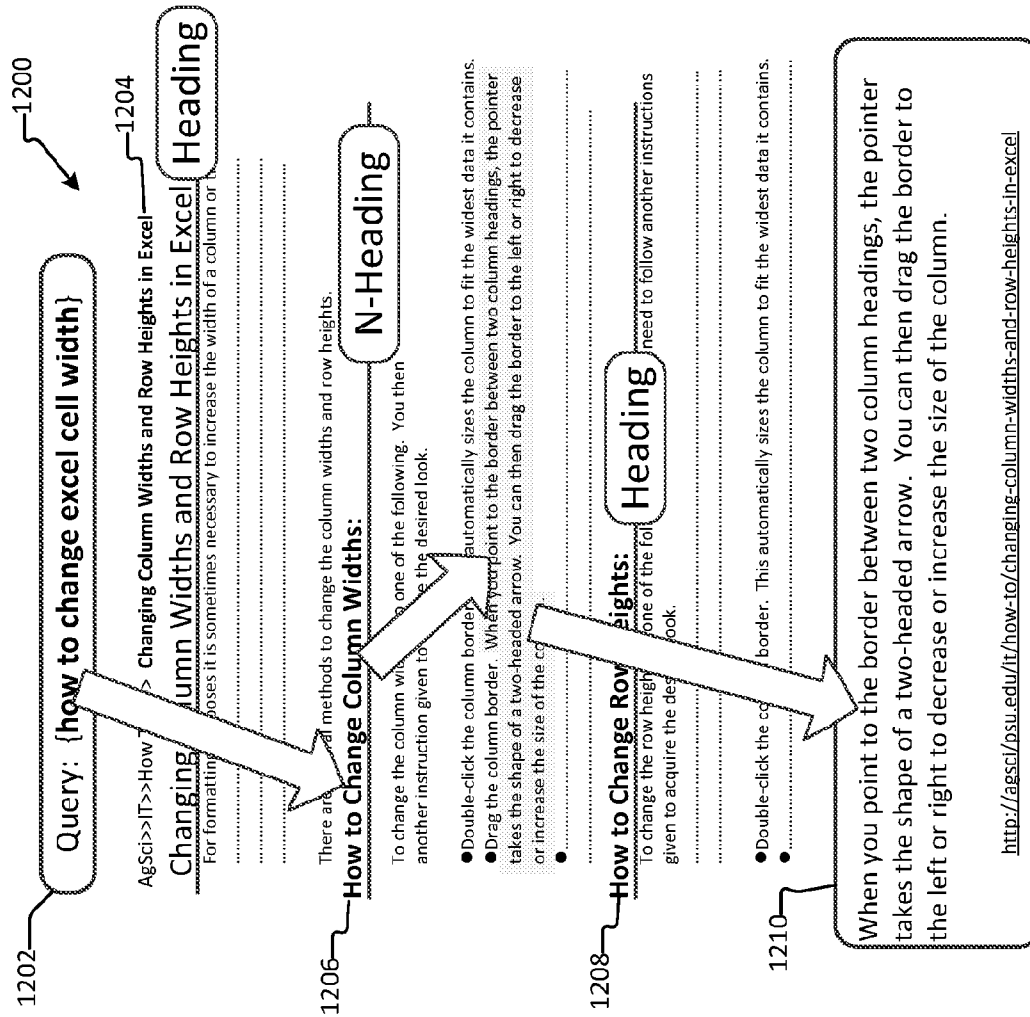


FIG. 12

1100

**Key Feature:**  
Contextual Matching

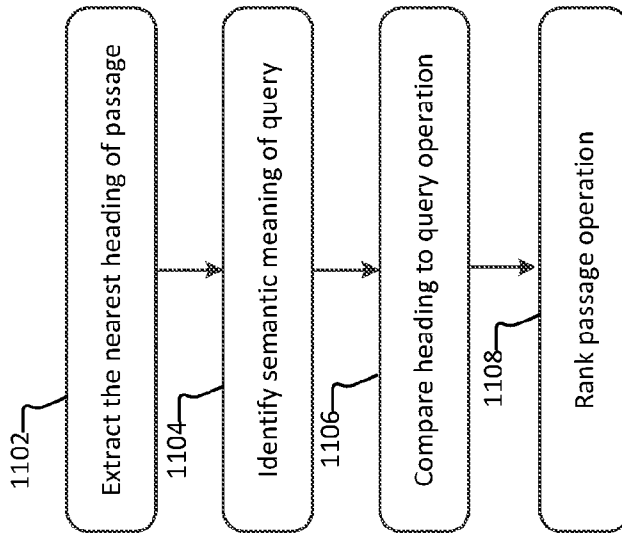


FIG. 11

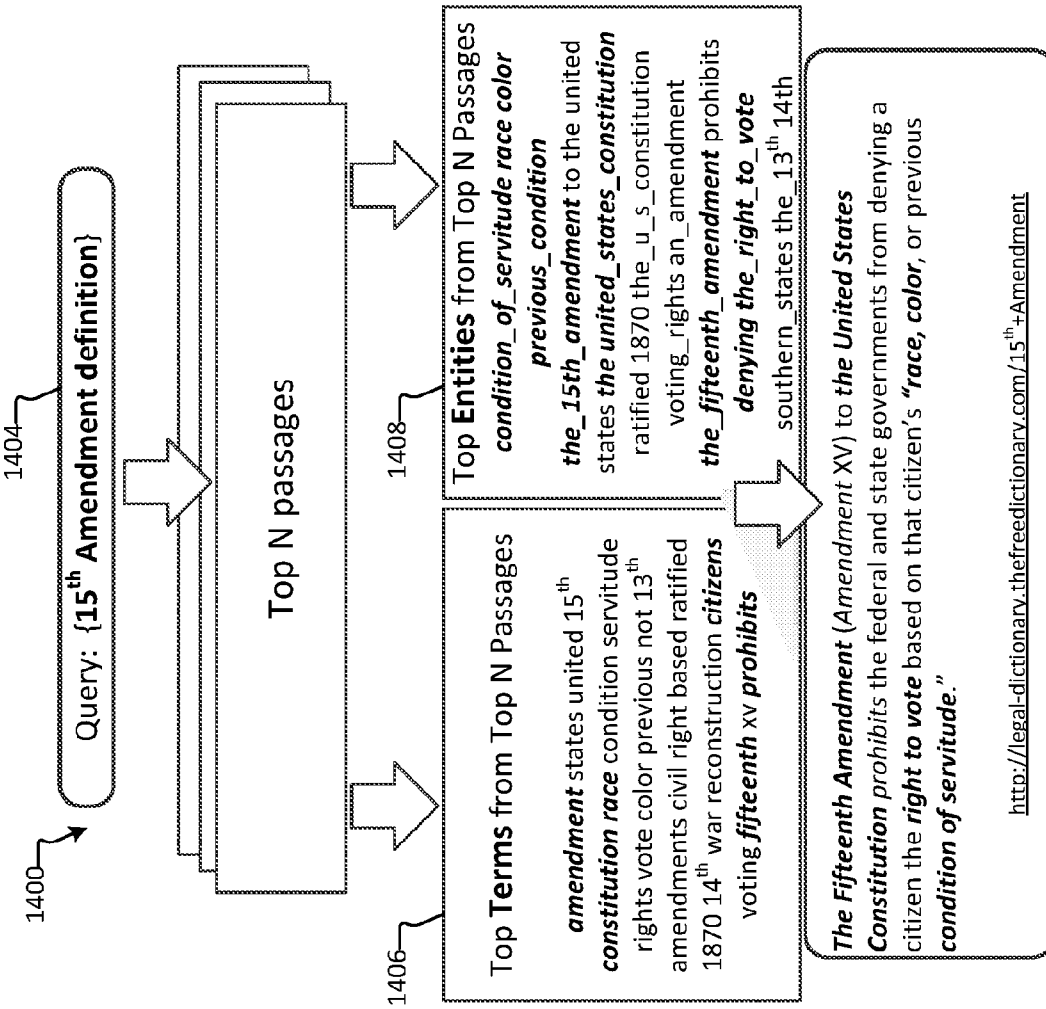


FIG. 14

Key Feature:  
 Re-ranking Service  
 Cross-document Aggregation




FIG. 13

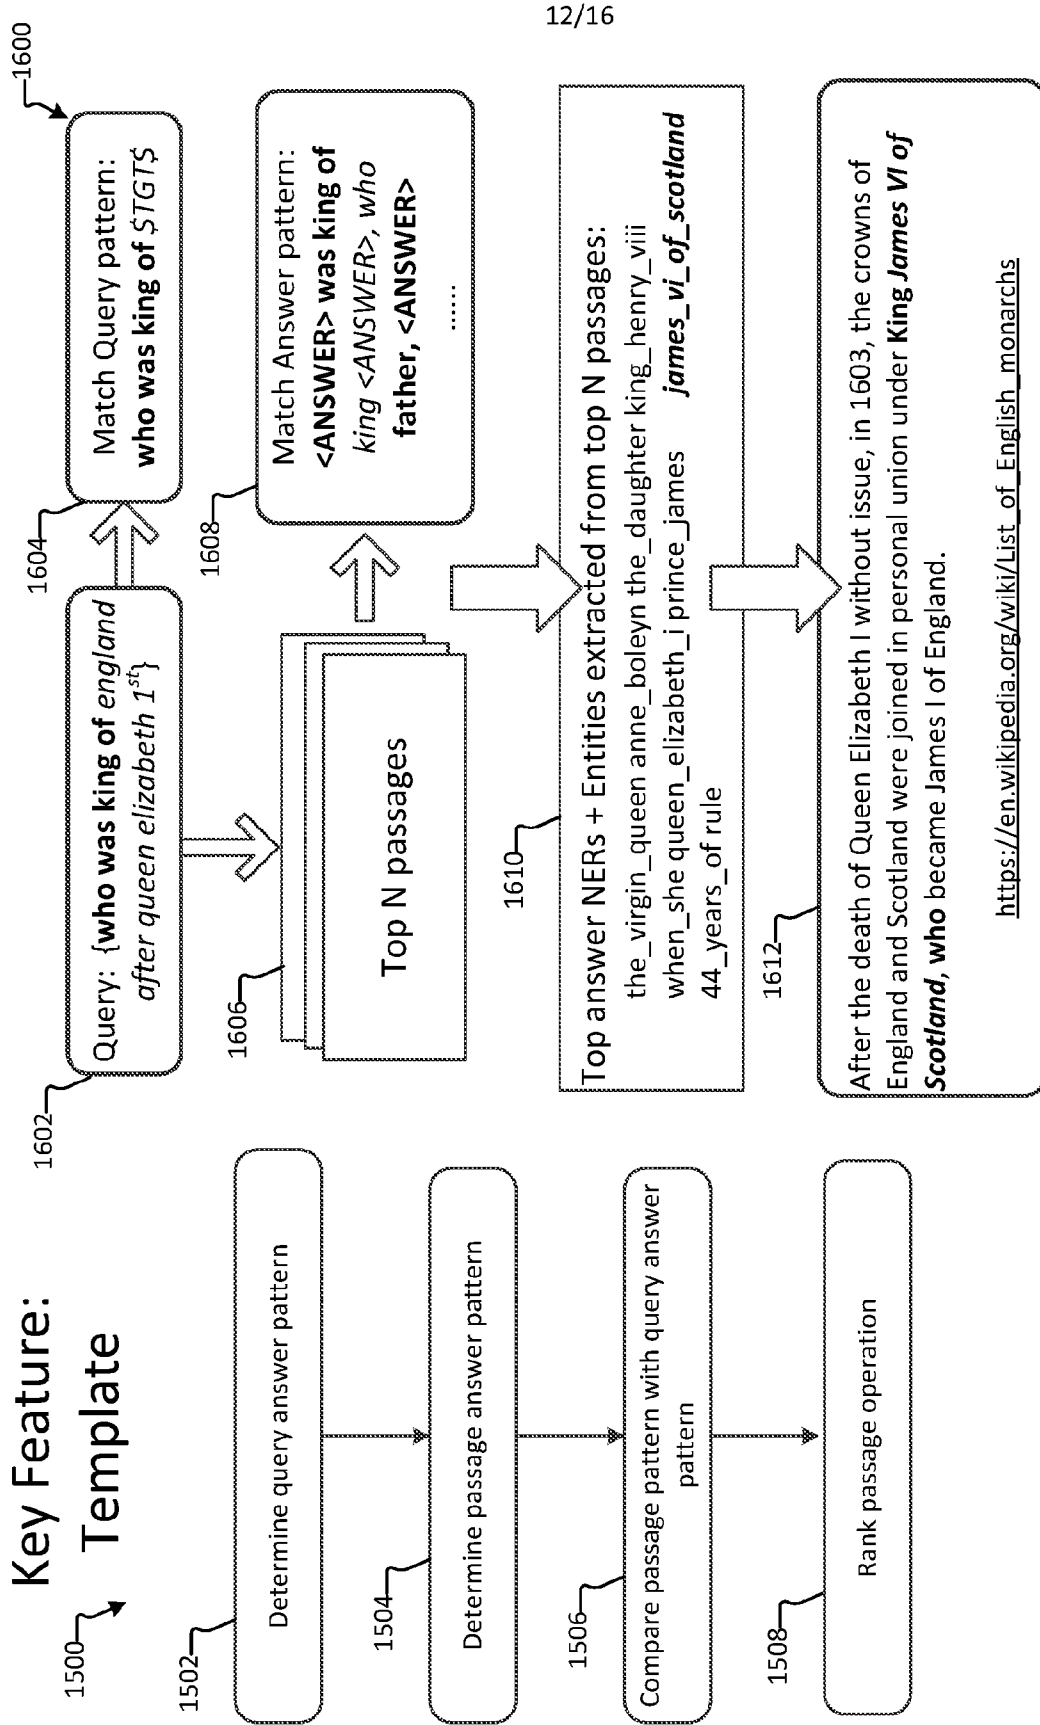


FIG. 15

FIG. 16

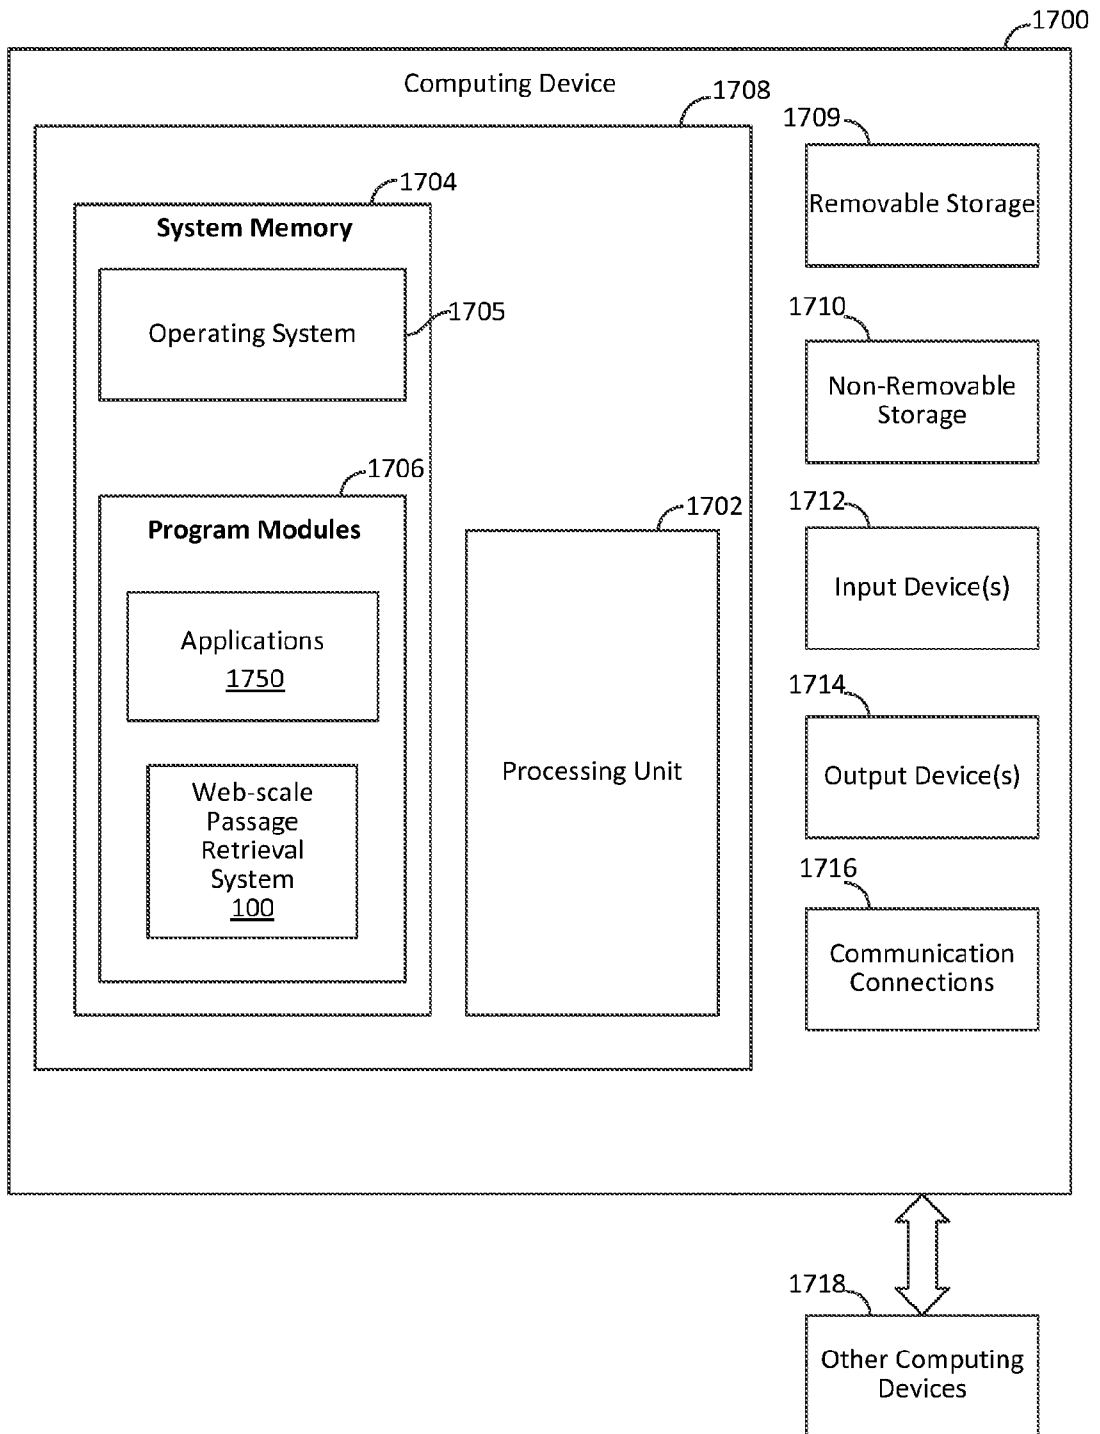



FIG. 17



MOBILE COMPUTING DEVICE

FIG. 18A

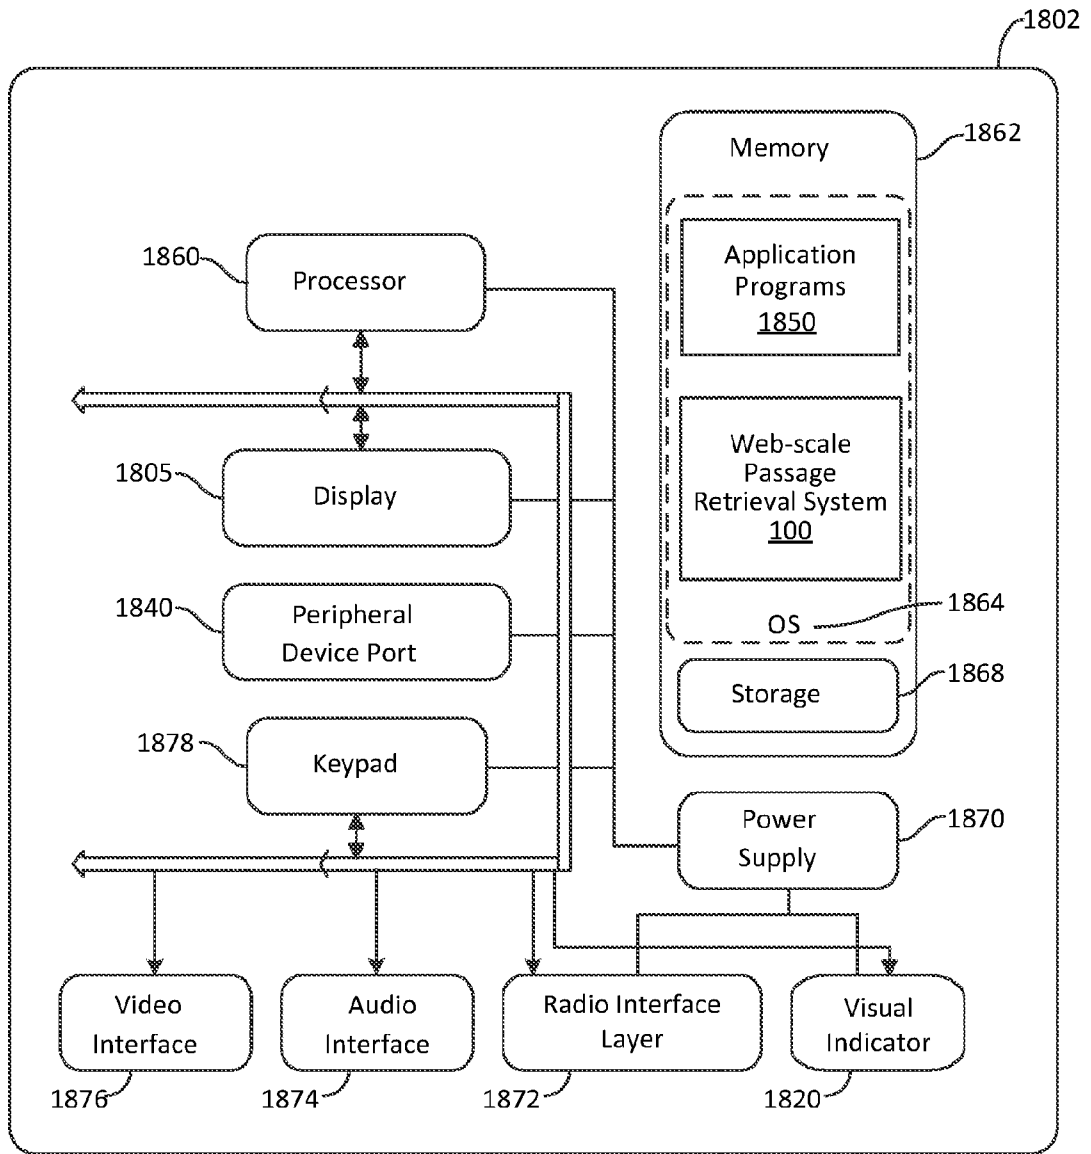


FIG. 18B

16/16

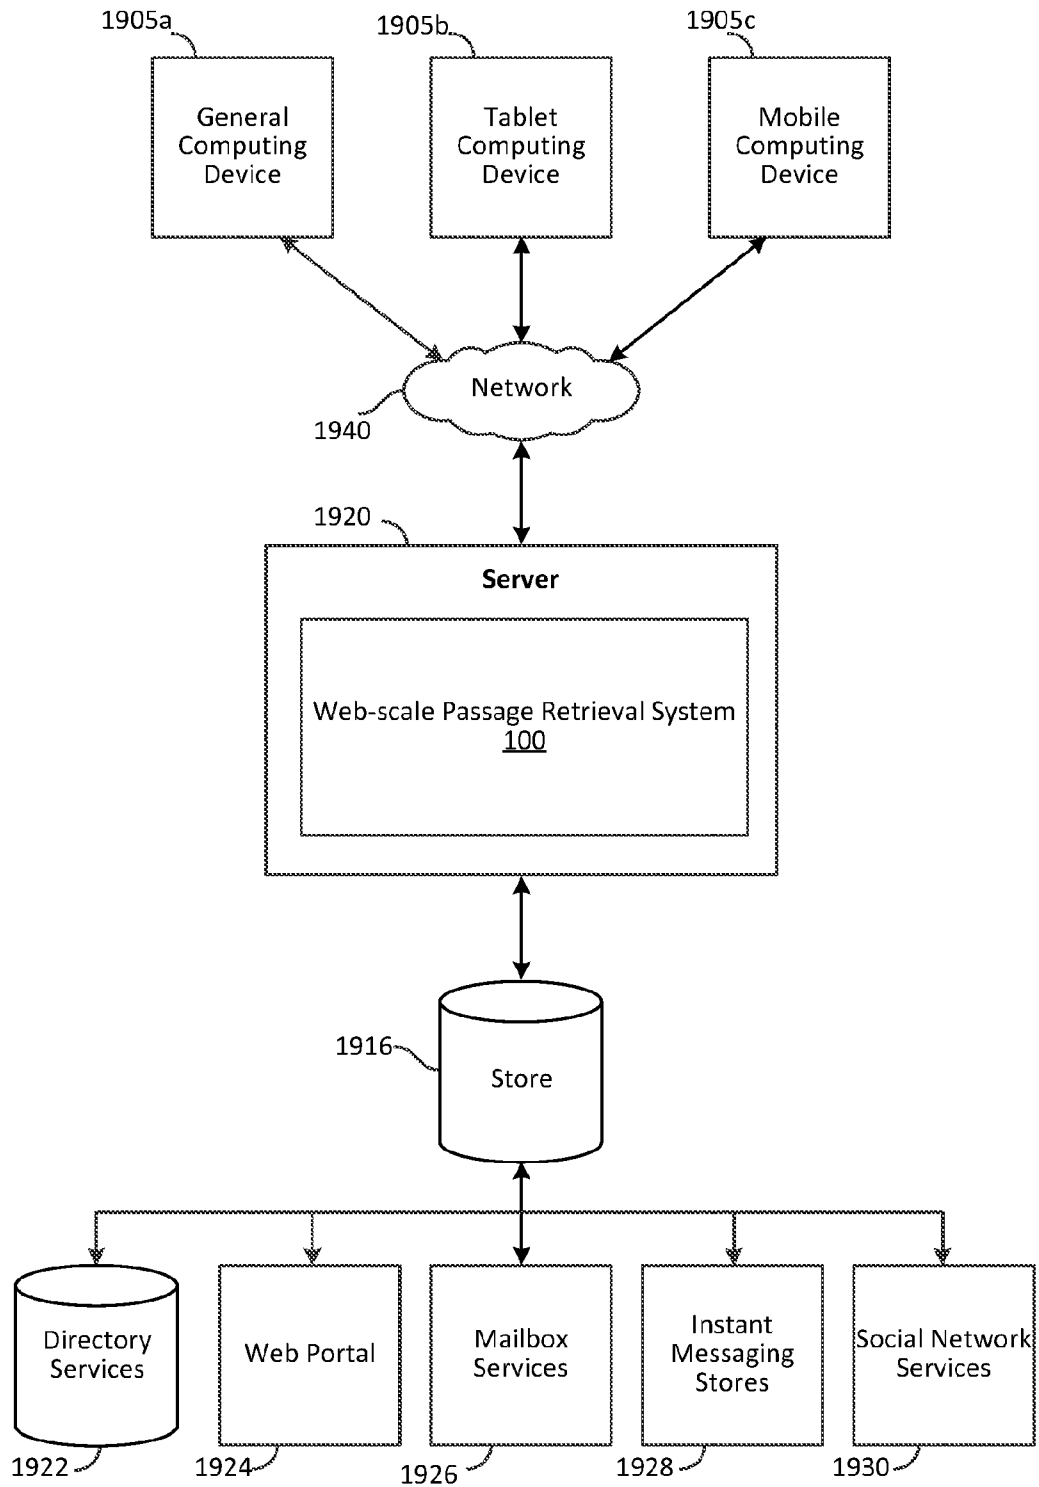


FIG. 19

**INTERNATIONAL SEARCH REPORT**

International application No.  
**PCT/CN2016/082980**

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>		
G06F 17/30(2006.01)i		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols)		
G06F		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
CNTXT;CNABS;DWPI;SIPOABS;CNKI: query, rank+, aggregat+, passage+		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2016015267 A1 (HEWLETT PACKARD DEVELOPMENT CO ET AL.) 04 February 2016 (2016-02-04) the whole document	1-20
A	WO 2009035871 A1 (POWERSSET INC) 19 March 2009 (2009-03-19) the whole document	1-20
A	US 2010185623 A1 (LU YUMAO ET AL.) 22 July 2010 (2010-07-22) the whole document	1-20
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search		Date of mailing of the international search report
<b>24 January 2017</b>		<b>24 February 2017</b>
Name and mailing address of the ISA/CN		Authorized officer
<b>STATE INTELLECTUAL PROPERTY OFFICE OF THE P.R.CHINA</b> <b>6, Xitucheng Rd., Jimen Bridge, Haidian District, Beijing 100088</b> <b>China</b>		<b>WANG, Yankun</b>
Facsimile No. <b>(86-10)62019451</b>		Telephone No. <b>(86-10)62089282</b>

**INTERNATIONAL SEARCH REPORT**  
**Information on patent family members**

International application No.

**PCT/CN2016/082980**

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)	Publication date (day/month/year)
WO 2016015267 A1	04 February 2016	None	
WO 2009035871 A1	19 March 2009	US 2009070322 A1	12 March 2009
US 2010185623 A1	22 July 2010	None	